Niche specificity, polygeny, and pleiotropy in herbivorous insects

[1]Nate B. Hardy and [2]Matt Forister

[1]Department of Entomology and Plant Pathology, Auburn University, Auburn, AL 36849

[2]Department of Biology, University of Nevada, Reno, NV 89557

Abstract

Why do herbivorous insects tend to be host specialists? Population genetic models predict specialization when there is antagonistic pleiotropy at a gene for host-use performance. But empirically, host-use performance is governed by many genetic regions, and antagonistic pleiotropy is rare. Here, we use individual-based quantitative genetic simulations to investigate the role of pleiotropy in the evolution of host-use specialization when host-use performance is polygenic. We find that if host-preference is allowed to evolve without cost, parasite populations tend to evolve host-use specialization even without pleiotropy; thus, it would seem that for a polygenic trait, the benefit of maintaining adaptive combinations of conditionally-neutral alleles suffices to drive specialization. But if there is a fecundity cost for host-preference, or if host patches are demographically volatile, host-use generalists evolve, even with high levels of pleiotropy. In sum, if pleiotropy is much more pervasive than has been observed in nature, our simulations show that it could play a role in driving the evolution of polygenic specialization. But pleiotropy is not necessary, and even when it is extensive, selection can favor generalist genotypes.

Introduction

Most species of herbivorous insects use as hosts only a small fraction of the plant species they encounter in their environment (Forister et al., 2015; Futuyma & Moreno, 1988), despite the obvious advantages of using many hosts and having a broader resource base (Hardy et al., 2020; Rainey et al., 2000; Smith & Hoekstra, 1980). How do we explain such host specificity? For the most part, theorists have assumed that host specificity is a consequence of antagonistic pleiotropy (Egas et al., 2004; Rueffler et al., 2004): an allele that confers high performance on one host confers low performance on another. Mathematical models show that if host-use performance is determined by one gene with antagonistic pleiotropy, and if populations can evolve host-use preferences, then the evolution of host-use specialization is all but inevitable (Ravigné et al., 2009), barring pronounced temporal heterogeneity in the host environment that could make specialization risky.

But in two key aspects, this theory bears little resemblance to what we know about the natural world. First, in herbivorous insects, most host-specific performance phenotypes, such as the ability to develop on a diet of leaves from a particular plant, seem to be controlled by many genes, not by one or even a few (Egan et al., 2015; Gompert et al., 2015). Second, host-associated alleles discovered in natural populations are often conditionally neutral, rather than antagonistically pleiotropic; they affect performance on one host but not another (Gompert et al., 2015). Conditionally neutral alleles might seem less likely to drive the evolution of host-use specialization. Moreover, it is not obvious what influence a few – or many – antagonistically pleiotropic loci might have in the evolution of a polygenic trait that is otherwise dominated by conditional neutrality. Here, we use forward-time individual-based quantitative genetic simulations to address this simple question: do simple genetic trade-offs matter in a polygenic world? More specifically, how might pleiotropy affect the evolution of specialization for a polygenic niche trait such as host-use in plant-parasitic insects?

Methods

Simulations were performed with SliM 3 (Haller & Messer, 2019). In this framework, models are specified in the Eidos programming language. Eidos codes are provided in S1-4. Before we describe our models in any detail, let us start with an overview. What we modeled was the mutational history of a population of diploid individuals that occurs in a simple environment consisting of two distinct resource patches. We imagined that these individuals were parasites, and that the resources patches were host species, called Host A and Host B. Individual host-use preference and performance phenotypes evolved as a function of mutations at many genetic loci. We were primarily interested in how changes in parasite genetic architecture – specifically, the ratio of pleiotropic to conditionally-neutral mutations – affected subpopulation divergence across host species. By subpopulation divergence we mean divergence at both neutral and host-use-affecting loci, with divergence at the latter sites reflecting the evolution of specialization.

Now let us look at the model details. As is typical for herbivorous insects, generations were discrete and non-overlapping. To keep things simple, individuals were hermaphroditic and mating was random. Selection was hard; the fitness, $\boldsymbol{W}$, of individuals was absolute (not relative to conspecifics). Moreover, $\boldsymbol{W}$ determined the odds of survival rather than fecundity. (How this was calculated is described below.) The regulation of each host-associated subpopulation was density-dependent, with a nominal carrying capacity for each host patch of 500 individuals. But since fitness was absolute, it was possible for a patch to sustain populations above or below that threshold. In other words, population size was an emergent property of the model, reflecting the balance between individual survival rates and fecundity (with individual survival probabilities, described in greater detail below, scaled by the ratio of local patch carrying capacity to the current number of individuals in the local subpopulation).

In brief, the parasite life cycle consists of a round of mate choice and reproduction (mutation, recombination, and offspring production), followed by a round of habitat choice, and then a round of

selection in which the survival probability of each individual was first determined by the match between their genotype and environment, and then scaled by subpopulation density (Figure 1). At the start of each generation, each surviving individual randomly selects as a mate one individual from their subpopulation. The fecundity of the choosing individual is determined by a draw from a Poisson distribution with a mean and variance (λ) of two. The selected number of crosses between a mating pair (chooser and chosen) are then added to the subpopulation. Note that choosers could also be chosen as mates by other individuals in their subpopulation (there is no limit to the number of matings an individual could be involved in, but the chances of being involved in more than one or two matings is small). Thus, individual fecundity could be much greater than that resulting from their own mate selection.

After reproduction, each new individual selects a host patch to inhabit. Habitat choice is determined in part by quantitative trait loci (QTLs). At the start of a simulation all individuals have no preference, that is, an equal likelihood of choosing Host A or Host B. Then, as the simulation progresses, host preference evolves as a function of the sum of mutational effects at preference-determining loci. The realized host choice behavior of each individual is calculated as a weighted average of their host-preference genotype and an inertial tendency to remain in their natal host patch. The weights were 40% for host preference and 60% for host inertia. Here are some more technical details. Host preferences phenotype values range from 0 to 1. At the start of a simulation, each individual has a preference genotype initialized at 0.5. Then, each generation, preference genotypes are updated as the sum of the initial genotype value (0.5) and the allele effects at preference loci. A separate variable for preference inertia is set to 0 on Host A and 1 on Host B. As an example, if the preference genotype value for an individual on Host B is 0.8 (0.5 + the sum of preference loci mutation effects), then overall preference becomes 0.4 * 0.8 + 0.6 * 1 = 0.92. When this weighted average is less than 0.5, Host A is preferred; when greater than 0.5 the other host is preferred. Preference is then

expressed by using the preference phenotype value as a cutoff for a draw from a random uniform distribution. To continue the same example, if the random draw were 0.6, the individual would choose Host B, and if the random draw were 0.99 the individual would choose and migrate to Host A, although such an outcome has a relatively low probability. There were no direct effects from host-preference alleles on survival, thus, selection on those loci was indirect via the fitness effects of alleles that determined performance on each host species (more on this later).

Now we turn to selection. On Host A, there is stabilizing selection on the quantitative phenotype value of Trait 1. One Host B, there is stabilizing selection on the phenotype value of Trait 2. As a concrete example, we can think of these traits as plant-defensive-chemical-detoxification systems – such as mixed-function oxidases and glutathione S-transferases (Enayati et al., 2005; Feyereisen, 1999) – with phenotype values corresponding to particular activity levels. Or the two traits could represent two components of a parasite's immune response to pathogens with different optimal values in different host-plant patches. Phenotypes are determined solely by genotypes, that is, by the additive combination of QTLs. We assume no environmental effects on phenotypes, or dominance effects on genotype values.

Each parasite individual has a diploid, 100 kb, one-chromosome genome. This genome is composed of ten 5 kb coding regions interspersed with ten 5 kb non-coding regions. Of course real herbivorous insect genomes are much larger and have multiple chromosomes. Working with smaller genomes is easier computationally, but increases linkage disequilibrium. To correct for that, we set the recombination rate to 1e-4 events per base pair (bp) per generation. At that rate, we expect about ten crossing-over events per genome per generation, which corresponds roughly to as much genetic shuffling as would be typical for a real herbivorous insect species with a five-chromosome genome (Wilfert et al., 2007). As alluded to previously, at the start of each simulation, all individuals have the same ur-genome. The mutation rate is 1e-7 per bp per generation. Five types of mutations could occur:

(1) neutral mutations in the non-coding regions, (2) pleiotropic mutations, that is, mutations at QTLs that govern Traits 1 and 2, and thereby affected survival on both Host A and B, (3) conditionally neutral mutations at QTLs that contribute to Trait 1 but not 2, (4) conditionally neutral mutations at QTLs that contribute to Trait 2 but not 1, and (5) pleiotropic mutations that governed the host-preference trait. Host-preference genes were mixed in with host-performance genes in the coding regions of the genome.

When a mutation occurred at a QTL affecting either preference or performance, an effect was drawn from a Gaussian distribution with a mean of zero and a variance $\boldsymbol{\sigma}$ which varied over simulations taking the values 0.1, 0.6, and 1.2. To be clear, the effect of an allele at a pleiotropic host-performance QTL is the same in both host patches (i.e. symmetrical), determined by one draw from the Gaussian distribution of effects. Whether the pleiotropy is negative or positive depends on the local patch optimum along with an individual's genetic environment, that is, the sum of allele affects across host-performance QTLs. Alternatively, we could have selected effects from a bivariate normal distribution with a covariance more than zero, but less than one. This would have allowed for less pronounced pleiotropies. We opted for maximally pronounced pleiotropies.

By the end of each simulation, the optimal performance trait value for each patch is displaced by ten units from zero, the starting condition (see below). Thus, when the allele effect variance is 1.2, the minimum number of averaged-sized mutations that it would take for a genotype value to match the final patch optimum is about eight, whereas when the allele effect variance is 0.1, the minimum is 100. But those are minimums. With each mutational effect drawn randomly from a Gaussian distribution, and without selection, it could take many more mutations to sum to a genotype absolute value of ten, and the disparity in the number of mutations needed when effect variances are small or large would be more pronounced. At the end of our simulations, the mean number of segregating alleles was 77.02,

with a standard deviation of 17.49. The mean number that reached a frequency greater than 90% in each host patch was 35.18, with a standard deviation of 27.71.

An individual's phenotype for each host-performance trait was determined by summing allele effects from pleiotropic and conditionally neutral loci. On Host A, an individual's fitness was determined by the distance between their genotype value for Trait 1, and the current Trait 1 optimum. Conversely on Host B, what mattered was the disparity between the Trait 2 genotype value and optimum. More specifically, fitness effects were determined with a Gaussian fitness function: We found the probability of observing a particular genotype value, given a Gaussian distribution with a mean equal to the trait optimum, and a variance $\boldsymbol{\omega}$ which sets the strength of selection (smaller variances equate to stronger selection). We then scaled that value in proportion to maximum probability in the Gaussian distribution, occurring at the mean, so that fitness ranged from 0 to 1. We ran our simulations with $\boldsymbol{\omega}$ set to 6.0, which can be interpreted as medium-strength selection.

Our main question was about how increasing the ratio of pleiotropic to conditionally neutral loci at host-performance determining sites – what we call the mutational mix – affected the evolution of host specialization, when host-use preference could also evolve. Thus, across each set of simulations, the mutational mix was set in turn to each of the following ratios: 2^{-10}, 2^{-5}, 2^{0}, 2^{5}, 2^{10} (from relatively few pleiotropic loci to many). We were also interested in how host-use evolution depended on (1) how host-patches diverged, (2) fecundity costs for strong host-use preference, and (3) the volatility of host-use patches.

Model 1. In the first model variant, we explored the case of a gradual host patch divergence and no-cost host-preference. At the start, the optimal host-use phenotype value in each host patch was the same, zero, which was also the starting mean genotype value for the parasites population. Then, after every 100 generations, the optimal quantitative phenotype value on Host A increases by a value of 0.2, while the optimal value on Host B decreased by the same amount. Starting host patch optima at zero

meant that initial parasite populations were well-adapted to their host environment. And the gradual divergence of host patch optima meant that host-use adaptation could proceed via incremental changes to genotype values. With no cost for strong host preferences, we expected that these conditions would be especially conducive to the evolution of host-use specialization. This and the other models are of course cartoons to help us think about the evolution of host use. Biologically, this first model could correspond to a situation in which environmental parameters (climate change) are gradually shifting the optima in different patches in different directions (perhaps associated with plants in contrasting microclimates). This scenario could also correspond to the gradual invasion of novel hosts that are adopted by different subpopulations; we have a variant of this model in which the host shifts in only one subpopulation, but we have presented the symmetrical case for simplicity.

Model 2. In this model variant, host patch divergence was gradual as in the first model, but we levied a fecundity cost on strong preference phenotypes, that is, we assumed there were trade-offs between primary metabolism and sophisticated host-sensing and host-seeking phenotypes (Hardy et al., 2020). We calculated a preference penalty as 1 / (1 + 0.1 * the sum of allele effects at preference QTLs); that penalty was then used to scale the λ for the Poisson distribution that determines fecundity (see above). We expected that such a cost would promote the evolution of un-choosy generalist populations, unless of course selection for generalism was opposed by antagonistic pleiotropy.

Model 3. In this model variant, there was a fecundity cost for strong host-preference phenotypes, as in Model 2, but host patch divergence was instantaneous. As in the gradual-divergence models, at the start of a simulation, the optimal phenotype value for both host patches was zero. Optima remained at zero for 1,000 generations, over which the strength of stabilizing selection was low, ω = 12.0. Then on the 1,001th generation, the optimum on Host A shifted to +10, the optimum on Host B shifted to -10, and the strength of selection increased dramatically to ω = 6.0. For the rest of the simulation, trait optima were fixed. Our aim was to have an initial phase in which populations

accumulated genetic diversity that was cryptically pleiotropic (McGuigan & Sgrò, 2009), that is, mutations that had no strong effect on fitness in the initial host environment, but that would have strong and antagonistic effects in the novel host environment. We though that this model captured genetic contingencies that could affect how a naive parasite population adapts to new hosts. We also suspected that a population reservoir of cryptically pleiotropic alleles would increase the overall impact of pleiotropy on host-use evolution.

Model 4. This model variant incorporated demographic stochasticity. Host-patch diverge was gradual, as in Models 1 and 2, and there was no direct fecundity cost on host-use preference phenotypes, as in Model 1. But in each patch, in each generation there was a 5% chance that the nominal patch carrying capacity would drop to 5% of its normal value. This model was inspired by the old idea that generalism is a kind of ecological bet-hedging (Jaenike, 1990). Patch volatility should disproportionately affect the fitness of specialist genotypes. A crash in a host patch almost wipes out a specialist subpopulation, whereas the same crash only shrinks a generalist population to ~55% of its previous state. As for Model 2, we expected the evolution of generalist populations, unless antagonistic pleiotropy got in the way.

We ran each simulation for 5,000 generations, with 100 replicates for each combination of model parameters. For different perspectives on the genetic states of a metapopulation at the end of each simulation, we logged these statistics: (1) $\boldsymbol{W}$: mean population fitness; (2) $\boldsymbol{F_{ST}}$: the standard ratio of between subpopulation neutral genetic variance to total neutral genetic variance; (3) the ratio of fixed pleiotropic to fixed conditionally neutral mutations, (4) the ratio of the mean effect sizes of fixed pleiotropic and conditionally neutral mutations, and (5) $\Delta\boldsymbol{W}$, a specialization index that measures subpopulation divergence, defined as the mean difference between individual fitness in the resident patch and what the fitness of an individual would have been in the other host patch, scaled by the

overall mean fitness of the metapopulation (with smaller values denoting more generalist populations because the individual-level discrepancy in fitness between patches is less).

In summary, we began by looking at a situation in which we thought the evolution of specialization was especially likely: where host-use phenotypes were rigid (without environmental effects), host-preference was cheap, the optimal phenotype value on one host species was distant from that on the other, and the host plant environment was demographically stable. We wanted to see if, when specialization was on order, pleiotropy played a role in delivering it. Then we considered two changes to this situation that we suspected would be more favorable to the evolution of generalists, namely, where host-preference had a cost, and the carrying capacities of host patches were volatile.

Results and Discussion

For each of the four model variants described above, we give five plots, each showing how a metapopulation variable tends to change with increasing amounts of pleiotropy in the mutational mix, as well as how those changes depend on, $\boldsymbol{\sigma}$, the mean allele effect size. First we show the mean fitness, $\boldsymbol{W}$, of individuals in each patch. In these plots, values near one denote near perfect matches between mean genotype values and local performance trait optima. To be clear, these plots just tell us how well-adapted the metapopulation is; they tell us nothing about the degree to which adaptation entailed divergence of host-use phenotypes between subpopulations. Next we show mean $\boldsymbol{F_{ST}}$ values. These can range between zero and one, with values closer to one denoting more pronounced neutral genetic divergence between subpopulations. In the third plot, we show mean values of the specialization index, $\Delta\boldsymbol{W}$. As a reminder, this also ranges between zero and one, with higher values corresponding to more pronounced divergence in host-use phenotypes between patches. In the fourth plot, we show the ratio of fixed pleiotropic to fixed conditionally neutral mutations, and in the fifth plot, we show the ratio of the mean effect sizes of the fixed pleiotropic and conditionally neutral mutations.

Model 1. No-cost host-preference in gradually-diverging host environments (**Figure 2**)

For this model, divergence between host patches was gradual. At the end of 5,000 generations, when $\boldsymbol{\sigma}$ (a variance parameter that controls the distribution of mutations on performance alleles) was 0.6 or 1.2 the mean fitness of the metapopulation was near one, in other words the population as a whole was well adapted. But when $\boldsymbol{\sigma}$ was 0.1, it was ~25% with lower with low-pleiotropy mutational mixes, and more than 90% lower with high levels of pleiotropy. How do we explain this? Well, as the average allele effect size decreases, the average number of sites needed to produce an optimal genotype value increases; decreasing $\boldsymbol{\sigma}$ may tend to increase the effective polygeny of a trait. With more polygeny, the challenge of generating sufficient genetic diversity and selecting for an optimal combination of alleles may become more challenging. Moreover, it would seem that this challenge becomes more significant with more pleiotropy

That story fits our other summary statistics. When $\boldsymbol{\sigma}$ was 0.1, $\boldsymbol{F_{ST}}$ and $\Delta\boldsymbol{W}$ (our index of host-use specificity) were near zero across the full range of mutational mixes; in other words, we ended up with poorly-adapted, genetically cohesive, generalist populations. When $\boldsymbol{\sigma}$ was 0.6 and pleiotropy was scant, $\boldsymbol{F_{ST}}$ and $\Delta\boldsymbol{W}$ were low, that is we ended up with well-adapted generalist parasite populations. Increasing pleiotropy increased population divergence and specialization only after the ratio of pleitropy in the mutational mix exceeded one. The picture was even more complicated when $\boldsymbol{\sigma}$ was 1.2, that is, for the largest allele effects that we considered. When pleitropy was either relatively rare, or much more frequent (>1000x) than conditional neutrality, $\boldsymbol{F_{ST}}$ and $\Delta\boldsymbol{W}$ were relatively high; we ended up with well-adapted, relatively-specialized subpopulations. But when the amount of pleiotropy in the mutation mix was medium-high (32x the rate of conditionally-neutral mutation), $\boldsymbol{F_{ST}}$ and $\Delta\boldsymbol{W}$ plummeted; thus, it seems that in some conditions increasing pleiotropy can actually decrease population divergence and specialization.

For further insight into the impact of pleiotropy we looked at two ratios: (1) the ratio of fixed pleiotropic to fixed conditionally neutral mutations, and (2) the ratio of the mean effect sizes of fixed pleiotropic and conditionally-neutral mutations. We saw that, when pleiotropic mutations are relatively rare, in comparison to conditionally neutral mutations, few pleiotropic mutations are fixed. Moreover, those that are fixed tend to have much smaller effects, especially when $\boldsymbol{\sigma}$ is not overly small. Simply put, selection purges large-effect pleiotropic mutations with extreme prejudice. It is only when pleiotropies are introduced to a population at a high rate that they are likely to exert influence on the adaptive fate of a population.

Models 2 and 3. Host-preference with a cost

In these models, strong host preference lowered individual fecundity. When the divergence of host-patch optima was gradual (Model 2**; Figure 3**), we found that such a cost was sufficient to evolve high-performance host-use generalists across a range of mutational mixes. As in the no-cost-preference simulations, when $\boldsymbol{\sigma}$ was low, so too was mean fitness, and that became more pronounced with increasing pleiotropy. Otherwise, fitness was near unity at the end of each simulation. The major changes caused by the addition of a performance cost were that $\boldsymbol{F_{ST}}$ and $\Delta\boldsymbol{W}$ were near zero, across the range of allele effect sizes, until the ratio of pleiotropic to conditionally neutral mutations was greater than 32-to-one – much higher than what has been observed for host-use traits in real plant-parasitic insect populations. If host-preference is expensive, parasite populations tend to evolve no preference or specialization.

What if host preference comes at a fecundity cost and the divergence in host-patch optima is instantaneous, after a long period of overlap and stasis? To find out, we ran simulations in which host-patches diverged all at once, and the strength of stabilizing selection increased dramatically after 1,000 generations. To reiterate, this lag could allow masked pleiotropic mutations to accumulate in a

population, and thus increase the potential impact of pleiotropy. Sure enough, this did appear to increase the range of mutational mix values over which pleiotropy increased parasite population divergence and specialization. Interestingly, it also dramatically reduced the mean effect sizes of fixed pleiotropic mutations. Thus, it seems that a naive population's mutational history can be a contingency that forces specialization in a newly diverse host environment. Even so, selection purges large effect pleiotropic alleles.

No-cost host-preference in fluctuating host patches (**Model 4; Figure 5**)

We also looked to see what would happen if host-preference evolves without a fecundity cost, but the carrying capacities of host patches were subject to stochastic crashes. As mentioned earlier, generalism should cut the fitness costs of such environmental volatility. What we found was similar to results from models in which host-use preference came with a fecundity cost. High-performance generalists evolved unless pleiotropic mutations were much more common than conditionally neutral mutations. The upshot is that we don't need direct costs for host preference to evolve polygenic generalism. The indirect cost on long-term fitness of specializing on a volatile resource is sufficient.

Caveats and roads not taken

Without a doubt, we have explored the parameter space in a willy-nilly fashion. For example, we could have reported on a range of selection strengths and recombination rates, and we could have had a more fully-factorial set of simulations for the parameters that we did vary. Such a thorough exploration might have produced a thoroughly unreadable tome. Instead, our goal has been more modest: to further the idea that the importance of antagonistic pleiotropy in the evolution of specialization continues to play too central of a role in the thinking of most empirical and theoretical ecologists (Fry, 1996). To that

end, we have attempted to point out a few neighborhoods in parameter space in which pleiotropy and or polygeny do or do not play roles in the evolution of host-use specificity.

In comparison to standard Wright-Fisher models, our models are complex. For example, we specified functions for mate choice, reproduction, habitat preference, and subpopulation-level density-dependent regulation. Many of the details of those functions were arbitrary, and future investigations could explore alternatives. On the other hand, in comparison to nature, our models are decidedly not complex. For example, our host environment only had two patches, we did not explicitly account for spatial relationships between host-patches or parasite individuals, and we did not consider how spatial and niche overlap could affect things like competition and mate choice. Moreover, we did not allow the environment to alter the mapping of host-use genotypes to phenotypes, despite much evidence of host-use phenotypic plasticity in nature.

Conclusions

If in herbivorous insect species the genetic architectures of traits affecting host-use performance were simple, pleiotropy could explain the preponderance of host-use specificity. But they are not simple; the traits affecting host-use performance tend to be polygenic. Of course, the genetic architecture of host use has been investigated in only the tiniest fraction of herbivorous insects. Nevertheless, our simulations, while leaving many neighborhoods of parameter space unexplored, have encompassed a fairly wide range of the architectural possibilities that we might encounter in nature. Across that range of simulated architectures, we find that with polygeny, pleiotropy is not a terribly powerful driver for the evolution of specificity unless it is much more prevalent than what has been observed in nature. Moreover, and perhaps more important given the history of theoretical work in the field, we also show that antagonistic pleiotropy is not needed to explain the predominance of specialized host associations. As long as host preference can itself evolve, then the evolution of specialization flows from polygeny

itself. The exception to that is if there are features of the genetic or external environment that add a cost to specialization. Then we see the evolution of the generalists. Thus we suggest that future experimental work on the evolution of specialization might not beat so loudly at the door of antagonistic pleiotropy (or the absence thereof) but should look more broadly for costs to generalism (Hardy et al., 2020).

Acknowledgments

Thanks to Zach Gompert, Josh Jahner, Gwendolyn Bird, and Noah Bevers for helpful comments on the manuscript. This work was supporting in part by NSF award 1744552 to NBH and YYY to MF.

Author contributions

MF came had the premise. NB and MF together developed the models and wrote the manuscript. NB ran the simulations.

References

Egan, S. P., Ragland, G. J., Assour, L., Powell, T. H. Q., Hood, G. R., Emrich, S., Nosil, P., & Feder, J. L. (2015). Experimental evidence of genome-wide impact of ecological selection during early stages of speciation-with-gene-flow. *Ecology Letters*, *18*(8), 817–825. https://doi.org/10.1111/ele.12460

Egas, M., Dieckmann, U., & Sabelis, M. W. (2004). Evolution restricts the coexistence of specialists and generalists: The role of trade-off structure. *American Naturalist*, *163*(4), 518–531. https://doi.org/10.1086/382599

Enayati, A. A., Ranson, H., & Hemingway, J. (2005). Insect glutathione transferases and insecticide resistance. In *Insect Molecular Biology* (Vol. 14, Issue 1, pp. 3–8). John Wiley & Sons, Ltd. https://doi.org/10.1111/j.1365-2583.2004.00529.x

Feyereisen, R. (1999). INSECT P450 ENZYMES. *Annual Review of Entomology*, *44*(1), 507–533. https://doi.org/10.1146/annurev.ento.44.1.507

Forister, M. L., Novotny, V., Panorska, A. K., Baje, L., Basset, Y., Butterill, P. T., Cizek, L., Coley, P. D., Dem, F., Diniz, I. R., Drozd, P., Fox, M., Glassmire, A. E., Hazen, R., Hrcek, J., Jahner, J. P., Kaman, O., Kozubowski, T. J., Kursar, T. A., … Dyer, L. A. (2015). The global distribution of diet breadth in insect herbivores. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(2), 442–447. https://doi.org/10.1073/pnas.1423042112

Fry, J. D. (1996). The Evolution of Host Specialization: Are Trade-Offs Overrated? *The American Naturalist*, *148*, S84–S107. https://doi.org/10.1086/285904

Futuyma, D. J., & Moreno, G. (1988). The Evolution of Ecological Specialization. *Annual Review of Ecology and Systematics*, *19*, 207–233.

Gompert, Z., Jahner, J. P., Scholl, C. F., Wilson, J. S., Lucas, L. K., Soria-Carrasco, V., Fordyce, J. A., Nice, C. C., Buerkle, C. A., & Forister, M. L. (2015). The evolution of novel host use is unlikely to be constrained by trade-offs or a lack of genetic variation. *Molecular Ecology*, *24*(11), 2777–2793. https://doi.org/10.1111/mec.13199

Haller, B. C., & Messer, P. W. (2019). SLiM 3: Forward Genetic Simulations Beyond the Wright-Fisher Model. *Molecular Biology and Evolution*, *36*(3), 632–637. https://doi.org/10.1093/molbev/msy228

Hardy, N. B., Kaczvinsky, C., Bird, G., & Normark, B. B. (2020). What We Don't Know About Diet-Breadth Evolution in Herbivorous Insects. *Annual Review of Ecology, Evolution, and Systematics*, *51*(1). https://doi.org/10.1146/annurev-ecolsys-011720-023322

Jaenike, J. (1990). Host Specialization in Phytophagous Insects. *Annual Review of Ecology and Systematics*, *21*, 243–273.

McGuigan, K., & Sgrò, C. M. (2009). Evolutionary consequences of cryptic genetic variation. *Trends in Ecology and Evolution*, *24*(6), 305–311. https://doi.org/10.1016/j.tree.2009.02.001

Rainey, P. B., Buckling, A., Kassen, R., & Travisano, M. (2000). The emergence and maintenance of diversity: insights from experimental bacterial populations. *Trends in Ecology & Evolution*, *15*(6), 243–247. https://doi.org/10.1016/S0169-5347(00)01871-1

Ravigné, V., Dieckmann, U., & Olivieri, I. (2009). Live where you thrive: joint evolution of habitat choice and local adaptation facilitates specialization and promotes diversity. *American Naturalist*, *174*(4). https://doi.org/10.1086/605369

Rueffler, C., Van Dooren, T. J. M., & Metz, J. A. J. (2004). Adaptive walks on changing landscapes: Levins' approach extended. *Theoretical Population Biology*, *65*(2), 165–178. https://doi.org/10.1016/j.tpb.2003.10.001

Smith, J. M., & Hoekstra, R. (1980). Polymorphism in a varied environment: how robust are the models? *Genetics Research*, *35*(01), 45. https://doi.org/10.1017/S0016672300013926

Wilfert, L., Gadau, J., & Schmid-Hempel, P. (2007). Variation in genomic recombination rates among animal taxa and the case of social insects. *Heredity*, *98*(4), 189–197. https://doi.org/10.1038/sj.hdy.6800950

Figures

Figure 1. Simulated parasite life cycle. Each generation consists of a round of random mate choice followed by reproduction, in which parent genomes are mutated and crossed, and then offspring are added to the occupied host patch. After reproducing, parents die. The environment consists of two host plant species. Offspring chose to stay on their natal host species, or migrate to the other host species. This decision is a function of an inertial tendency to stay where they are, and their host preference genotypes, which are determined by the sum of allele effects at many QTLs. Selection determines survival probabilities; survival is more likely for individuals that have quantitative host-use genotypes that more closely match the optimum value on their host plant species.

Figure 2. Model 1. No-cost host-preference in gradually-diverging host environments. Coral-colored lines are for mean allele effect sizes, $\boldsymbol{\sigma}$, of 1.2, green lines are for $\boldsymbol{\sigma}$ of 0.6, and blue lines are for $\boldsymbol{\sigma}$ of 0.1. (A) Mean fitness, $\boldsymbol{W}$, of individuals across parasite metapopulation. Values near one denote near perfect matches between mean genotype values and local performance trait optima. (B) Mean $\boldsymbol{F_{ST}}$ values. Values closer to one denoting more pronounced neutral genetic divergence between subpopulations. (C) Mean values of the specialization index, $\Delta\boldsymbol{W}$**,** values closer to one denote more pronounced divergence in host-use phenotypes between patches. (D) Ratio of fixed plieotropic to fixed conditionally neutral mutations. (E) Ratio of the mean effect sizes of the fixed pleiotropic and conditionally neutral mutations.

Figure 3. Model 2. Expensive host-preference in gradually-diverging host environments. Coral-colored lines are for mean allele effect sizes, $\boldsymbol{\sigma}$, of 1.2, green lines are for $\boldsymbol{\sigma}$ of 0.6, and blue lines are for $\boldsymbol{\sigma}$ of 0.1. (A) Mean fitness, $\boldsymbol{W}$, of individuals across parasite metapopulation. Values near one denote near perfect matches between mean genotype values and local performance trait optima. (B) Mean $\boldsymbol{F_{ST}}$ values. Values closer to one denoting more pronounced neutral genetic divergence between subpopulations. (C) Mean values of the specialization index, $\Delta\boldsymbol{W}$**,** values closer to one denote more pronounced divergence in host-use phenotypes between patches. (D) Ratio of fixed plieotropic to fixed conditionally neutral mutations. (E) Ratio of the mean effect sizes of the fixed pleiotropic and conditionally neutral mutations.

Figure 4. Model 3. Expensive host-preference in instantaneously-diverging host environments. Coral-colored lines are for mean allele effect sizes, $\boldsymbol{\sigma}$, of 1.2, green lines are for $\boldsymbol{\sigma}$ of 0.6, and blue lines are for $\boldsymbol{\sigma}$ of 0.1. (A) Mean fitness, $\boldsymbol{W}$, of individuals across parasite metapopulation. Values near one denote near perfect matches between mean genotype values and local performance trait optima. (B) Mean $\boldsymbol{F_{ST}}$ values. Values closer to one denoting more pronounced neutral genetic divergence between subpopulations. (C) Mean values of the specialization index, $\Delta\boldsymbol{W}$**,** values closer to one denote more pronounced divergence in host-use phenotypes between patches. (D) Ratio of fixed plieotropic to fixed conditionally neutral mutations. (E) Ratio of the mean effect sizes of the fixed pleiotropic and conditionally neutral mutations.

Figure 5. Model 4. No-cost host-preference in demographically-volatile host environments. Coral-colored lines are for mean allele effect sizes, $\boldsymbol{\sigma}$, of 1.2, green lines are for $\boldsymbol{\sigma}$ of 0.6, and blue lines are for $\boldsymbol{\sigma}$ of 0.1. (A) Mean fitness, $\boldsymbol{W}$, of individuals across parasite metapopulation. Values near one denote

near perfect matches between mean genotype values and local performance trait optima. (B) Mean $\boldsymbol{F_{ST}}$ values. Values closer to one denoting more pronounced neutral genetic divergence between subpopulations. (C) Mean values of the specialization index, $\Delta \boldsymbol{W}$**,** values closer to one denote more pronounced divergence in host-use phenotypes between patches. (D) Ratio of fixed plieotropic to fixed conditionally neutral mutations. (E) Ratio of the mean effect sizes of the fixed pleiotropic and conditionally neutral mutations.

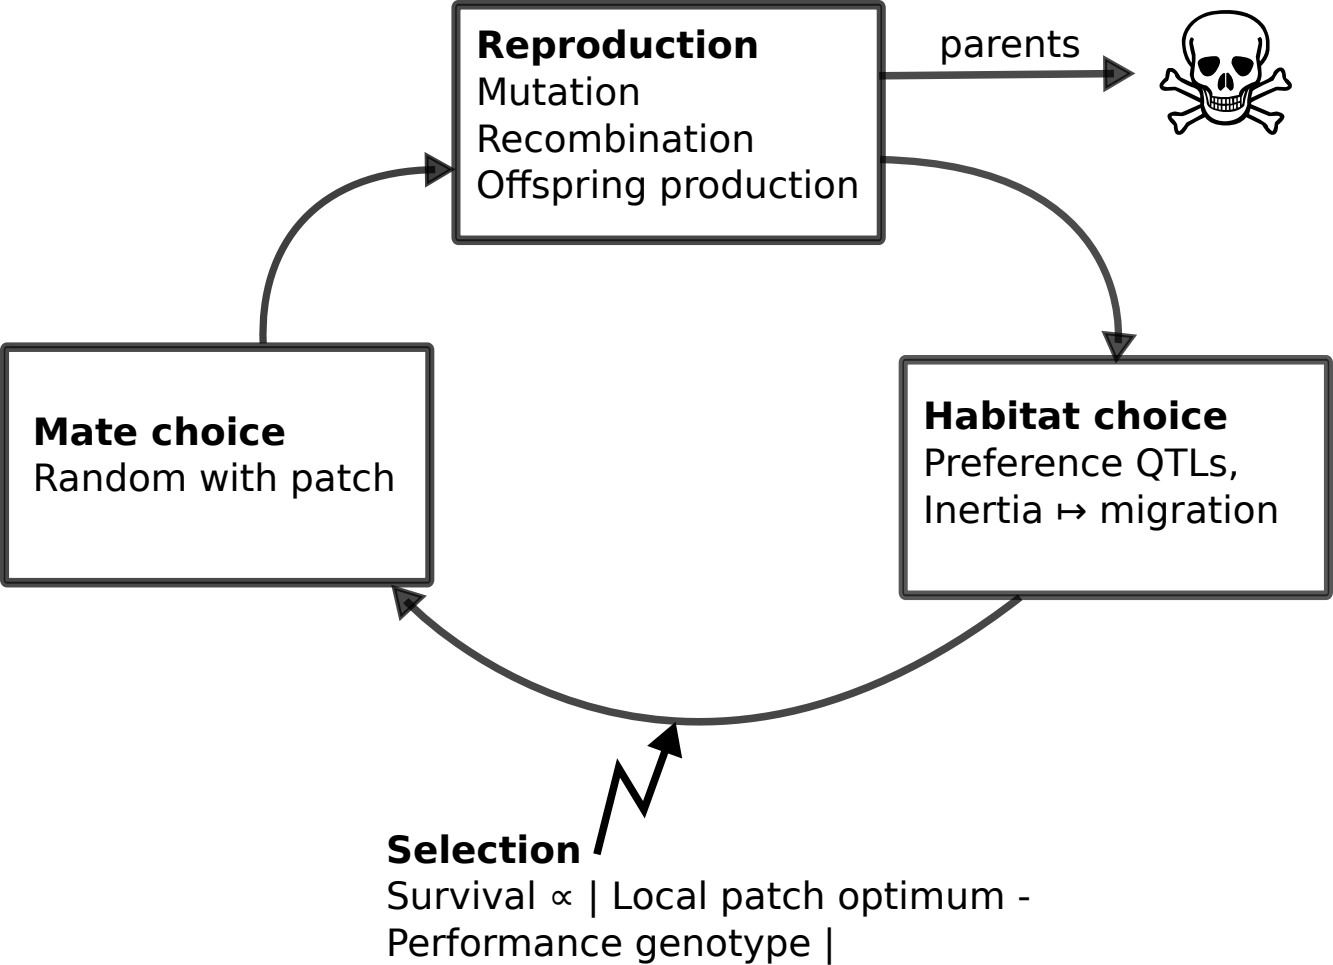
Reproduction
Mutation
Recombination
Offspring production
parents
Habitat choice
Preference QTLs,
Inertia ↦ migration
Mate choice
Random with patch
Selection
Survival ∝ | Local patch optimum -
Performance genotype |

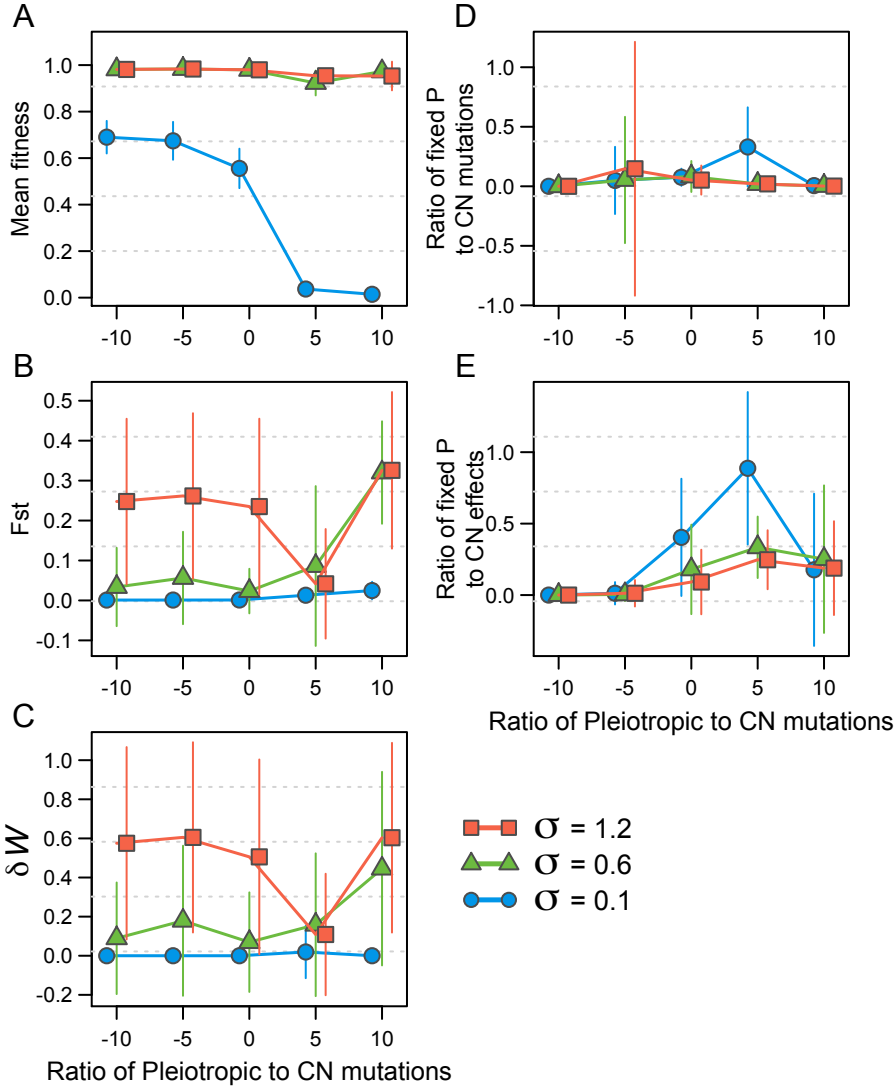
A
Mean fitness
1.0
0.8
0.6
0.4
0.2
0.0
-10
-5
0
5
10
B
Fst
0.5
0.4
0.3
0.2
0.1
0.0
-0.1
C
δW
1.0
0.8
0.6
0.4
0.2
0.0
-0.2
Ratio of Pleiotropic to CN mutations
D
Ratio of fixed P to CN mutations
1.0
0.5
0.0
-0.5
-1.0
E
Ratio of fixed P to CN effects
1.0
0.5
0.0
Ratio of Pleiotropic to CN mutations
σ = 1.2
σ = 0.6
σ = 0.1

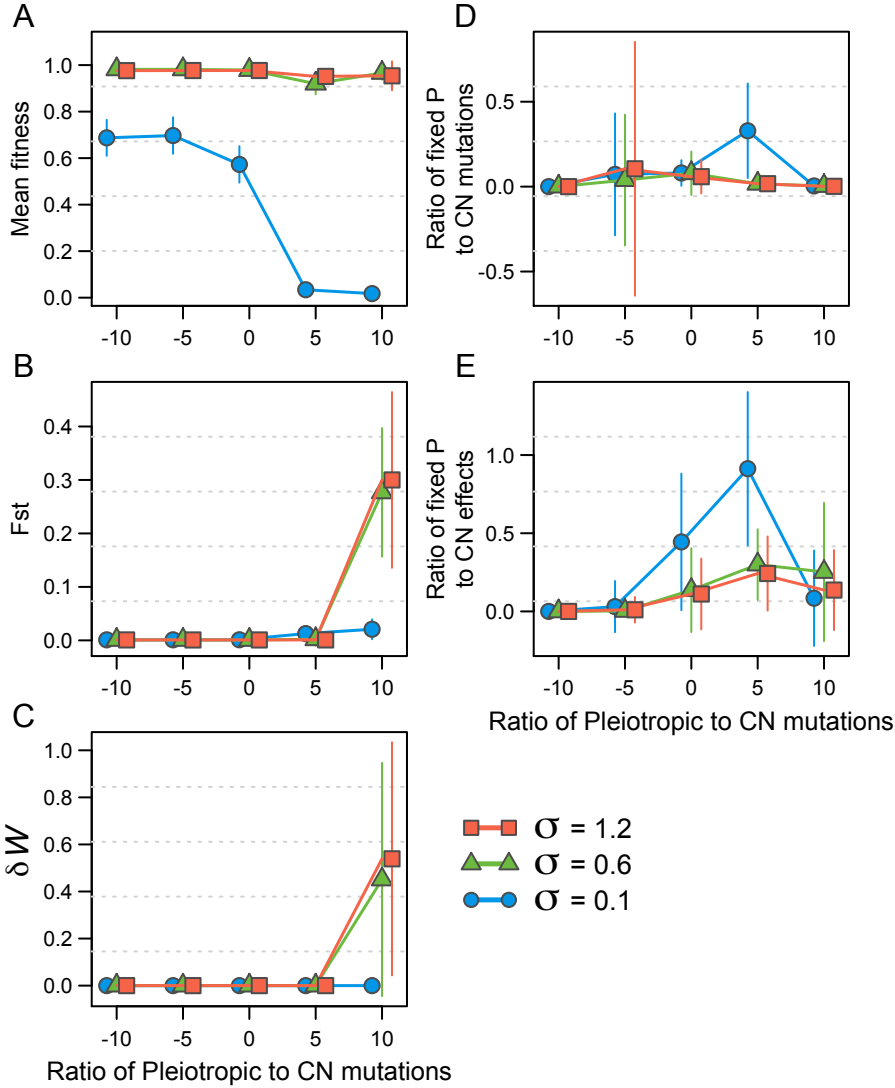
A
Mean fitness
1.0
0.8
0.6
0.4
0.2
0.0
-10
-5
0
5
10
B
Fst
0.4
0.3
0.2
0.1
0.0
C
δW
Ratio of Pleiotropic to CN mutations
D
Ratio of fixed P to CN mutations
0.5
0.0
-0.5
E
Ratio of fixed P to CN effects
1.0
0.5
0.0
Ratio of Pleiotropic to CN mutations
σ = 1.2
σ = 0.6
σ = 0.1

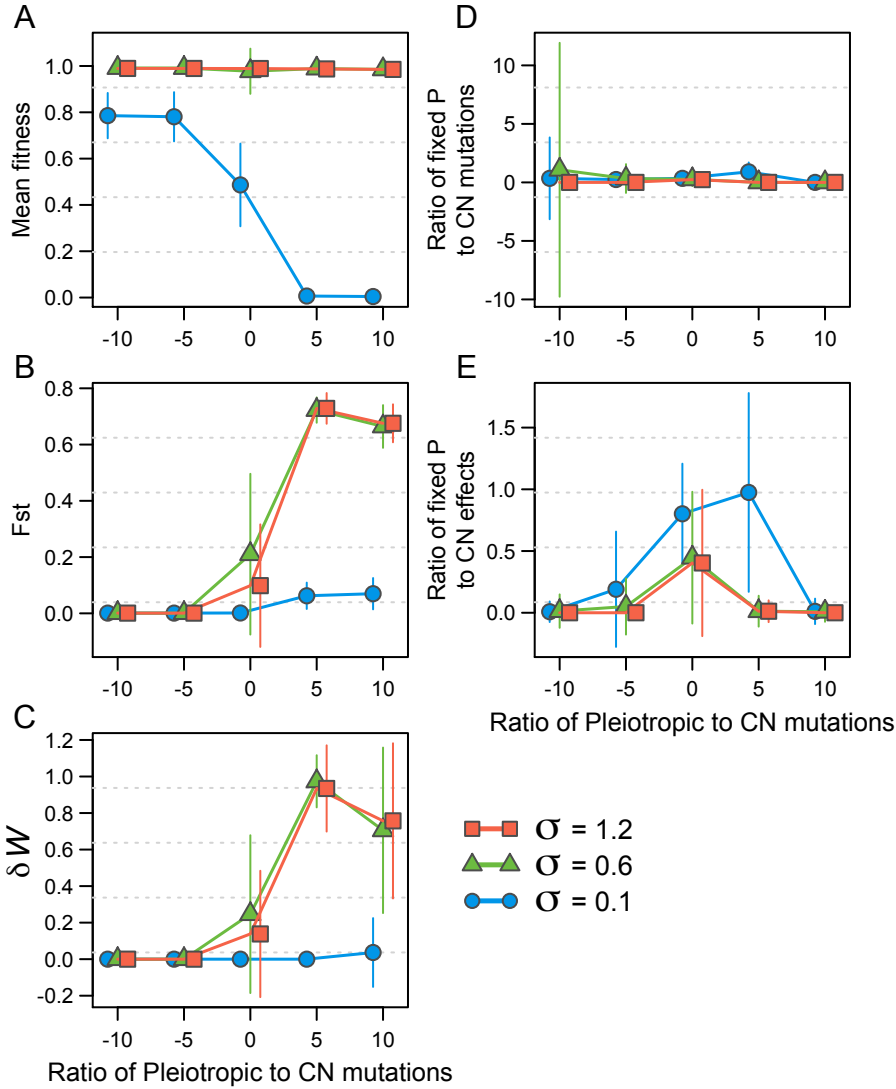
A
Mean fitness
1.0
0.8
0.6
0.4
0.2
0.0
-10
-5
0
5
10
B
Fst
0.8
0.6
0.4
0.2
0.0
C
δW
1.2
1.0
0.8
0.6
0.4
0.2
0.0
-0.2
Ratio of Pleiotropic to CN mutations
D
Ratio of fixed P to CN mutations
10
5
0
-5
-10
E
Ratio of fixed P to CN effects
1.5
1.0
0.5
0.0
Ratio of Pleiotropic to CN mutations
σ = 1.2
σ = 0.6
σ = 0.1

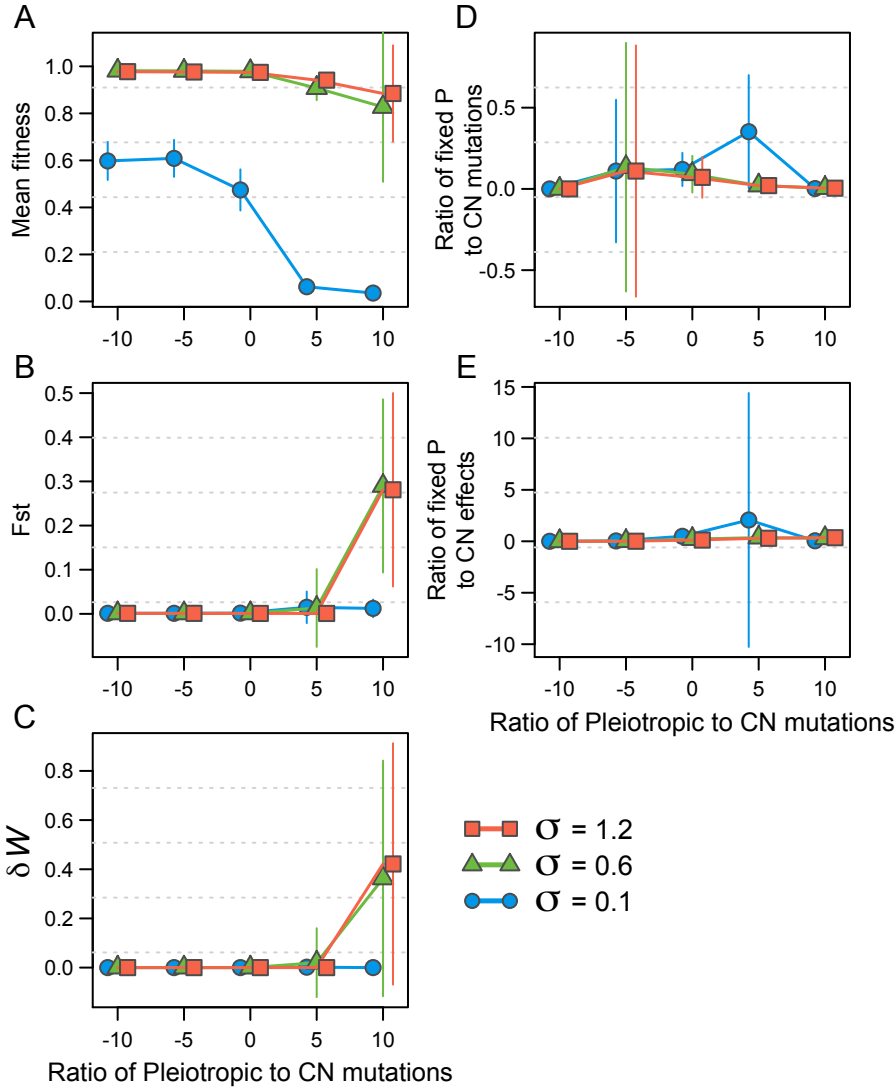
A
Mean fitness
1.0
0.8
0.6
0.4
0.2
0.0
-10
-5
0
5
10
B
Fst
0.5
0.4
0.3
0.2
0.1
0.0
C
δ W
0.8
0.6
0.4
0.2
0.0
Ratio of Pleiotropic to CN mutations
D
Ratio of fixed P to CN mutations
0.5
0.0
-0.5
E
Ratio of fixed P to CN effects
15
10
5
0
-5
-10
Ratio of Pleiotropic to CN mutations
σ = 1.2
σ = 0.6
σ = 0.1